Journal of the American Society for Information Science

Interfaces for Distributed Systems of Information Servers

Au: Page makeup rules will be followed before your article is printed.
Ed.

Brewster Kahle and Harry Morris
WAIS Inc., 1040 Noel Drive, Menlo Park, CA 94025

Jonathan Goldman
Thinking Machines Corporation, 1010 El Camino Real, Menlo Park, CA 94025

Thomas Erickson
Apple Computer, 20525 Mariani Avenue, MS 76-3H, Cupertino, CA 95014

John Curran
NEARnet, 10 Moulton Street, Cambridge, MA 02138

Au: Please cite Figure 1 through Figure 11 in text.

Interfaces for information access and retrieval are a long way from the ideal of the electronic book that you can cuddle up with in bed. Nevertheless, today's interfaces are coming closer to supporting browsing, selection, and retrieval of remote information by nontechnical users.

This article describes five interfaces to distributed systems of servers that have been designed and implemented: WAIStation for the Macintosh, XWAIS for X-Windows, GWAIS for Gnu-Emacs, SWAIS for dumb terminals, and Rosebud for the Macintosh. These interfaces talk to one of two server systems: the Wide Area Information Server (WAIS) system on the internet, and the Rosebud Server System, on an internal network at Apple Computer. Both server systems are built on Z39.50, a standard protocol, and thus support access to a wide range of remote databases.

The interfaces described here reflect a variety of design constraints. Such constraints range from the mundane—coping with dumb terminals and limited screen space—to the challenging. Among the challenges addressed are how to provide passive alerts, how to make information easily scannable, and how to support retrieval and browsing by nontechnical users. There are a variety of other issues which have received little or no attention, including budgeting money for access to "for pay" databases, privacy, and thus how to assist users in finding out which of a large (changing) set of databases holds relevant information. We hope that the challenges we have identified, as well as the existence and public availability of source code for the WAIS system, will serve as a stimulus for further design work on interfaces for information retrieval.

Introduction

It requires little prescience to predict that one day computers will put an ocean of information at the finger tips of a vast population of users. However, although there is a considerable amount of information available from remote sources, the bulk of it is accessible only to information professionals, or users with technical backgrounds. A variety of obstacles effectively block the ordinary user from accessing information via the computer. These obstacles include the difficulty of locating appropriate information sources, the cumbersome maneuvers needed to get online and to connect to remote sources, and cryptic query languages. Furthermore, even if a user has succeeded in accessing a remote information source, it is likely that it will have its own special purpose interface, which may or may not support the user's needs.

In this article, we describe two systems—Wide Area Information Severs (WAIS) and Rosebud—which provide a protocol-based mechanism for accessing a variety of remote, full-text information servers. These systems have the potential for supporting a single interface to a wide variety of information sources, and offer a good platform on which to explore the design of interfaces for information retrieval. After a summary of existing information retrieval systems, we describe the server systems, and then describe the five interfaces to them. In the course of these descriptions we discuss design constraints, interface issues, and practical maters which had an impact on the designs. We conclude with a summary, and some remarks on important issues that have not been addressed, and a invitation for other investigators to use the WAIS system as a platform for exploring interfaces to multiple, remote information sources.

 CCC 0002-8231/93/000000-00

Background

Existing Systems

While a review of all existing systems is beyond the scope of this article, it is useful to list a number of the most popular or significant interfaces for information retrieval.

Commercial interfaces for accessing full-text resources on computers can be broken down into dialup services, local file access, and LAN-based access tools. Dialup systems such as Dialog and Dow Jones offer TTY interfaces to users, with menus and command lines being the dominant access tools. Some dialup services are offering client programs that run on personal computers to add graphical interfaces such as "Navigator" by Compuserve. In general, these interfaces are unique to the information provider. Local file access through full-text indexing has been achieved in command line form (e.g., the Unix command "grep") and in screen-based interfaces (e.g., ON Location [ON], and Digital Librarian [NeXT]). These interfaces often give browsing and searching capabilities for local files. Some of these interfaces have been stretched to work with files on file servers. LAN-based access tools usually use some sort of query language to access servers on the net, such as Verity's Topic system (VERITY), and numerous library systems. These query languages require some user training. Integrated tools for cross-platform, cross-vendor information access are not currently available in other systems.

A variety of research projects have explored information retrieval systems. The SuperBook project (Egan, 1989) targets users of static information. Project Mercury (Ginther-Webster, 1990) is a remote library searching system that uses a client-server model. Information Lens (Malone, 1986) is a structured e-mail system for assisting in managing corporate information. NetLib for software (Dongarra, 1987) and Mosis for information on how to fabricate chips (Mosis) are examples of e-mail–based information retrieval systems.

The WAIS and Rosebud Projects

The two systems of information servers described in this article grew out of two, partially entwined projects: WAIS and Rosebud. A goal of both projects was to define an open protocol that would allow any user interface or information server that talked to the protocol to interact with any other component that used the protocol. From the user's perspective, this would mean that user interfaces and information sources could be mixed and matched, according to the user's needs.

WAIS started as a joint project between Thinking Machines Corporation, Apple Computer, Dow Jones & Co., and KPMG Peat Marwick (Kahle & Medlar, 1991). The proximate goal was to define the open protocol and demonstrate its feasibility by implementing and demonstrating a multivendor system which provided ordinary users with access to a variety of remote databases. Thinking Machines contributed its Connection Machine–based retrieval technology, Apple contributed its expertise in user studies and interface design, and Dow Jones & Co. provided access to its commercial information sources. KPMG Peat Marwick provided access to its corporate data, and served as a site for user studies and testing. The WAIS system was installed KPMG Peat Marwick and enabled the designers to study the success of the system in a real world context. The WAIS system uses pseudo–natural language queries, relevance feedback to refine queries, and accesses full-text, unstructured information sources. These technologies were used because they had already been tested independently, thereby leading to faster implementation of the complete system. The WAIS system will be described in more detail in the next section.

During the same period, the Rosebud project was underway within Apple. Rosebud's goal was to serve as an internal platform for research into system architecture and human interface issues and, as a consequence, employed a variety of more experimental technologies and was tested in-house. Like WAIS, Rosebud was based on user studies conducted at KPMG Peat Marwick, and used the same underlying protocol, Z39.50. The details of the Rosebud Server System will be described in a different study.

After the collaborative phase of the WAIS project came the Internet experiment. In this phase of WAIS, source code for the open protocol, information servers, and several interfaces were made freely available over the Internet. In addition, Thinking Machines established and maintained a directory of information servers that WAIS users could query to find out about available information sources. This phase of WAIS is still in progress, and has resulted in the creation of new interfaces, the availability over the Internet of more than 100 servers on three continents, and over 100,000 searches of the directory of servers. In the first six months of the Internet experiment, approximately 4000 users from 20 countries have tried this system, with no training other than documentation (Kahle, Goldman, Morris, & Shen, 1991). Administrators of popular information servers indicate that they are getting over 50 accesses a day from many countries.

The WAIS System

WAIS employs a client-server model using a standard protocol (based on Z39.50) to allow users to find and

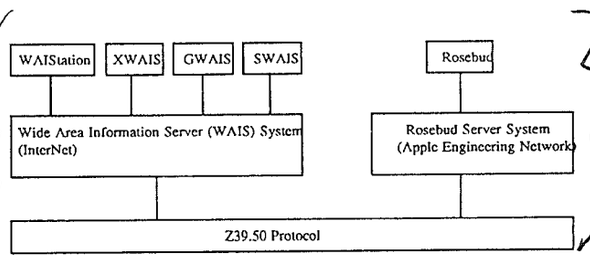

FIG. 1. The interfaces to the WAIS and Rosebud server systems, and the protocol.

retrieve information from a large number of servers. The client program is the user interface, the server does the indexing and retrieval of documents, and the protocol is used to transmit the queries and responses. Any client which is capable of translating a user's request into the standard protocol can be used in the system. Likewise, any server capable of answering a request encoded in the protocol can be used.

A WAIS server can be located anywhere that one's workstation has access: on the local machine, on a network, or on the other end of a modem. The user's workstation keeps track of a variety of information about each server. The public information about a server includes how to contact it, a description of the contents, and the access cost.

The WAIS protocol (Davis et al., 1990) is an extension of the existing Z39.50 standard (NISO, 1988) from NISO. It has been augmented where necessary to incorporate many of the needs of a full-text information retrieval system. To allow future flexibility, the standard does not restrict the query language or the data format of the information to be retrieved. Nonetheless, a query convention has been established for the existing servers and clients. The resulting WAIS protocol is general enough to be implemented on a variety of communications systems.

The WAIS clients will be described in detail in the next several sections. However, all of them work in a basically similar way. On the client side, queries are expressed as strings of words, often pseudo–natural language questions. The client application then packages the query in the WAIS protocol and transmits it over a network to one or more servers. The servers receive the transmission, translate the received packet into their own query languages, and search for documents satisfying the query. The lists of relevant documents are then encoded in the protocol and transmitted back to the client. The client decodes the response and displays the results. The documents can then be retrieved from the server. The documents can be in any format that the client can display such as word processor files or pictures.

WAIStation: An Interactive Query Interface

WAIStation at a Glance

Target Machine	Macintosh Plus and above, 9" Monochrome screen
Effort	1 person-year
Number of Users	2000
Status	Finished, freely distributed
Language	ThinkC
Communications	TCP/IP and Modem (not supported)
Designer	Harry Morris
Organization	Thinking Machines
Availability	Available for anonymous FTP from /public/wais/WAIStation*.sit.hqx @think.com
Design Goals	Implementable quickly, support interactive queries well, changeable based on user's comments, make something very simple to learn (partner friendly). Try out many ideas: interactive queries, passive alerting, asking multiple servers.
Used	In a study with accountants and tax consultants at KPMG: very good user acceptance. In the Internet experiment: estimated that half of the uses of WAIS are using WAIStation. (Based on when the directory of servers did not work for Macintoshes, usage dropped to half).
Problems	Dealing with the directory of servers (s). Modem code was difficult to get right.

WAIStation was designed for use in the WAIS experiment at KPMG Peat Marwick. As such, we needed an interface that would be easy to use, and would encourage successful searches by users untrained in search techniques. Peat Marwick often sends its employees into the field toting their Macintosh SEs along for use as portable computers. Thus we had to design the interface to run on a nine-inch black-and-white screen, and make minimal demands on CPU and memory. Furthermore, WAIStation was designed for use over modems and slow LANs.

Design Rationale

In designing WAIStation, we were informed by two metaphors—search as conversation and storage by file folder. The process of formulating an effective search is highly interactive. Of the documents which match a query, the ones which match "best" are displayed. One or more may be of interest, in which case, they can be fed back to the system, interactively improving the search. We choose to view this process as a conversation. Thus the initial natural language question becomes ~~that~~ the starting point for give and take between the user and the server(s). Relevance feedback provides the context for the question. As the search proceeds, some results may suggest alternative searches or branches of the conversation. This is provided for by allowing several questions to evolve at the same time.

Eventually, one or more questions may be refined to the point where they are finding consistently good results. At this point, the question can be automated, becoming a dynamically updated file folder. At intervals these questions wake up and query their servers. The results are stored in the results field for later inspection. They can be thought of as regular Macintosh folders, except as augmented with a charter describing how to keep their contents up to date.

Ed: *Note* Author asked to provide this single spaced c in paragraphs (col 1)

This parallel with the Macintosh folder structure suggested a drag and drop construction for the user interface itself. Constructing a question is a three-step process—typing the key words, specifying the servers to use, and specifying the relevant documents of feedback. If we think of questions like Macintosh folders, we can use the Macintosh's drag-and-drop mechanism for putting sources and relevant documents into a question. This approach makes WAIStation's mechanics instantly familiar to users of the Macintosh finder.

Fig 2, 3, 4 + 5 should be moved within the WAIStation section.

Human Interface

When WAIStation starts up, two windows appear—one contains the users available Sources (see ~~below~~ Figure 2) and one contains the users saved Questions. Sources are identified by an eye icon, questions by a question mark icon.

Double clicking on a question icon opens the stored question, including any new results found since the last time it was examined. The top half of the question window contains a field in which to type key words (the natural language part of the question), a list of relevant documents, a list of sources, and a list of result headlines. Sources can be added to the question by selecting a source icon (in the Sources window) and dragging it into the question. Relevant documents are specified in the same way.

Result documents, returned by the servers, can be examined by double clicking on their icon (see Figure 3 + 4). Note that the result list contains a graphical indication of how well each document matches the query. The original graphic was a series of zero to four stars, similar to the ratings found in *TV Guide*. We thought that this rating scheme would be easily recognized. Experience proved that the stars did not provide enough information to be recognized or to discriminate among the documents. Latter versions of the software replaced the stars with a horizontal bar giving 20 levels of resolution.

Any of the resulting documents can be opened and viewed in its own window. WAIStation supports plain ASCII documents as well as PICT format pictures. Text windows automatically scroll to the position which the server considered the most relevant part of the document. This allows the user to quickly determine if a file is useful. In order to perform well over slow communications channels (modems and slow LANs) the text is downloaded on demand in 15 line chunks. The keywords used in the query are automatically highlighted in bold.

Sources are specially formatted text files which describe information servers and how to get to them. Double clicking on a source displays a window with several controls. The top part is information specified by the server itself: (see Fig 5)

- a pop-up menu to specify the method of contacting the server (ip-address/tcp-port, modem number and speed, or location of a local index);
- a script to run after logging in (for use by modems);
- a database to search (servers can support multiple databases);
- a display of when the server is updated, how much it costs to search; and
- a textual description of the databases' contents.

The bottom half of the source window allows the user to specify personal information about the server:

- when to contact it (for automatic update);
- when it was last contacted;
- how much to spend on it;
- how much credence its results should be given (this is used to scale document scores, which helps in the sorting of responses to questions asked of multiple servers);
- the number of documents to ask for when searching it; and
- the font and type size to use when displaying plain text results (important to publishers).

Several of these fields are merely placeholders in the current implementation. In particular, budget and confidence have not been implemented yet since there are no for-pay servers yet, and the number of sources is still relatively small.

Source files can also be retrieved from servers. This allows users to search ~~serves~~ servers whose database elements are pointers to other servers. The results can be used as targets for further searches. An experimental directory of ~~severs~~ servers is being maintained on the Internet.

Implementation

WAIStation was implemented in ThinkC 4.0 using the object oriented class library. It took about a man-year of effort. The most difficult parts were the automatic update facility and the communications. Automatic update required the ability to do background processing—which is not a normal part of the Macintosh operating system. Communications were difficult primarily because we were simultaneously debugging the Z39.50 protocol, modem code, and the (then new) Apple Communications Toolbox. We eventually left modems unsupported, and replaced the Communications Toolbox with direct calls to MacTCP. Through this experience we found that communications speeds of less than 9600 baud were barely tolerable for interactive text retrieval.

Observations

We estimate the WAIStation is now in use by over 2000 users in 20 countries. The common user complaints center around configuring MacTCP, using (the undocumented) directory-of-servers, and avoiding a bug requiring the software to be installed on the start-up disk.

We have noticed several shortcomings in the current design:

- Users want access to their own data. WAIStation is capable of searching a Macintosh-based inverted index file, but we unbundled the index builder when we realized how much work it would take to make it useful under Macintosh OS. OnLocation (On Technology) is an implementation of a Macintosh indexer that could be used.
- Interaction with the directory of servers is incomplete. It is not obvious which search results are source files, and what to do with the ones that are. It should be possible to drag a retrieved source directly into a question's source window, but the present interface requires that it be saved first. The lesson we learned was that special cases should be handled specially, rather than forcing users to use general techniques "for consistency's sake."
- Printing documents and searching for keywords in documents (find/find-next) are simple functions which users expect.

- People want to see their documents in their original form. WAIStation currently only displays ASCII and PICT. This can be fixed with format filters such as Claris' XTND, at the expense of the ability to download arbitrary sections of a document, since such filters require that the document be processed from the beginning.
- Relevance feedback was not obvious. Users unfamiliar with the use of relevance feedback did not think to use it—it needs to be made more automatic. One way to do this might be to extend the notion that a question is a conversation, with relevance feedback as context (or body language)—clients or servers can be written that watch their users, and deduce which documents were relevant based on which ones were read. A simpler approach might be to always do relevance feedback, presenting the results in a "see also" list. We tried this, but the Macintosh was too slow to make it useful.
- Communications over 2400-baud modems are too slow to support interactive queries. We found that 9600 baud is barely acceptable, while 56 Kb is sufficient to support several users.
- The finder-like interface (drag and drop) is not obvious. Even though the Macintosh finder is based on drag and drop, no one expected it in an application. Once users were shown what to do, it became natural. It was also not necessarily the best use of screen space, since it required that both the start and end of the drag be visible on the screen at the same time. Another anomaly worth mentioning is that, although we were simulating the finder, we had no "trash can" analogy. Removing a source was accomplished by dragging it onto the desk top and dropping it there, which confused some users.
- The alerting system was crude. For example, there was no visual cue to tell the user that a question had found new documents in the background. Also, the background searches did not exclude previously read documents.
- Headlines often do not give enough context. The headlines displayed in the question window were only about 60 characters long, making it difficult to identify which documents were useful without opening them. Furthermore, there was no provision to display the document's date or the name of the source it came from.

X-Windows Based Interface for WAIS: XWAIS

Ed: *Note* Author has been asked to supply single spacing ē in paragraphs

XWAIS at a Glance

Target Machine	X-Windows terminals on UNIX machines
Effort	4 person-months
Number of Users	500
Status	Finished, freely distributed
Language	C
Communications	TCP/IP
Designer	Jonathan Goldman
Organization	Thinking Machines
Availability	Available anonymous FTP from /public/wais/wais*.tar.Z@think.com
Design Goals	Copy WAIStation so that we can leverage one design, portable, and based only on freeware, Display data in many different formats (image, text, etc.).
Used	Used in the Internet experiment. Heavy use by X users within Thinking Machines and outside.
Problems	Installing it has caused many users to stumble. The number of variables (architectures, X directory structures) makes it difficult to make it portable, touch on the ability to handle different types (this is unique to this interface). Uses other programs to help (like interapplication communication).

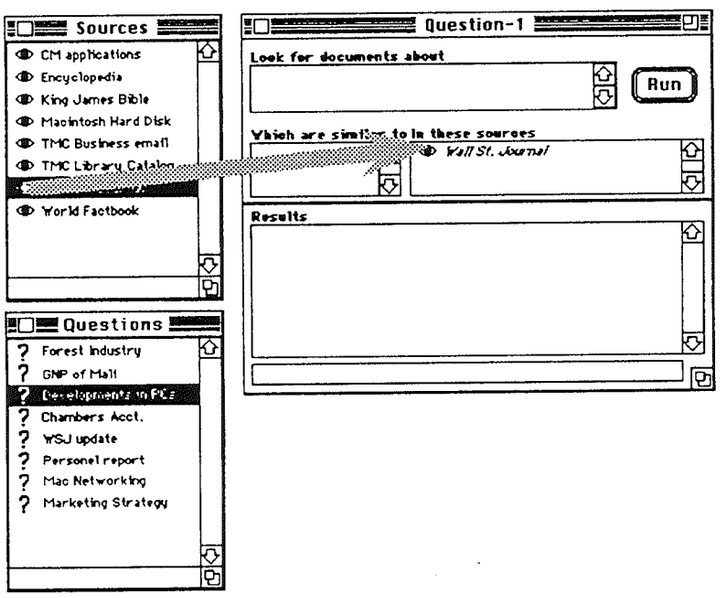

FIG. 2. WAIStation's Sources and Questions windows store the user's personal objects. Dragging a source into a question window specifies that the question will contact the source in order to fulfill its charter.

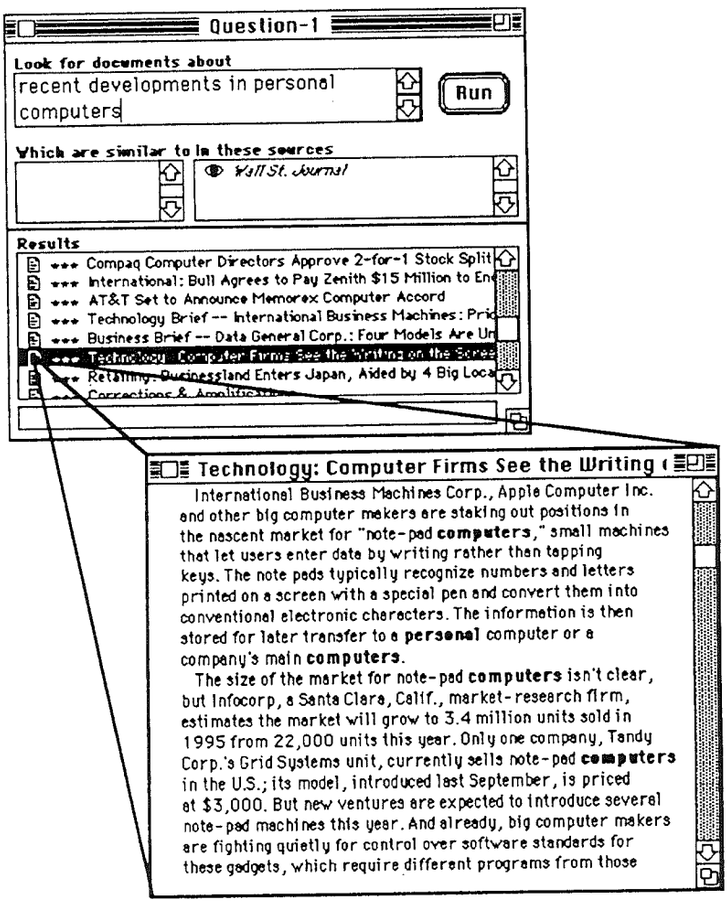

FIG. 3. After running the question, results are displayed in a scrolling list. Double clicking on a result opens a document window. Query words are highlighted.

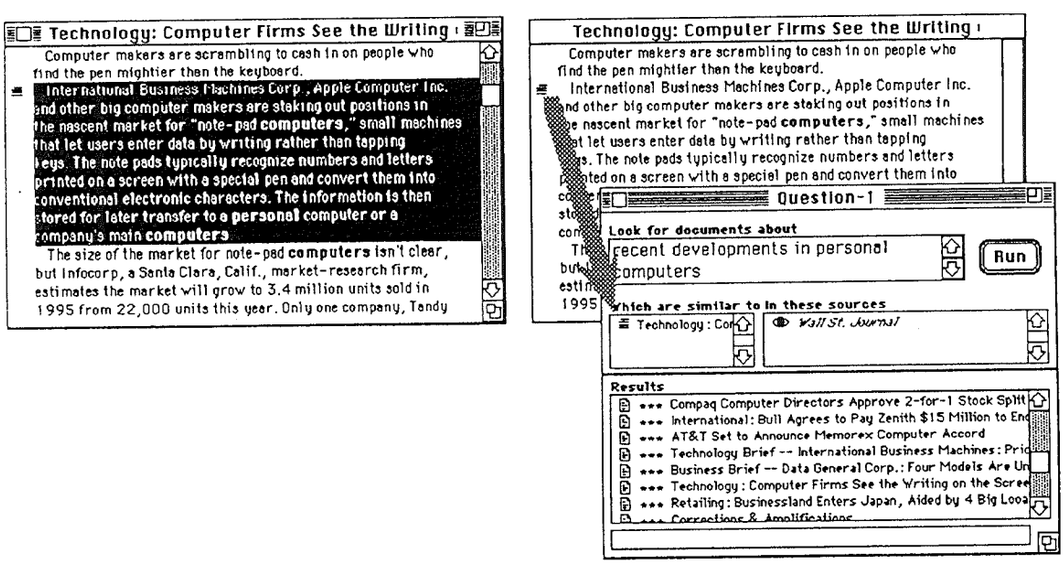

FIG. 4. Relevance feedback is done by selecting a document or part of a document and dragging the document or paragraph icon into a question.

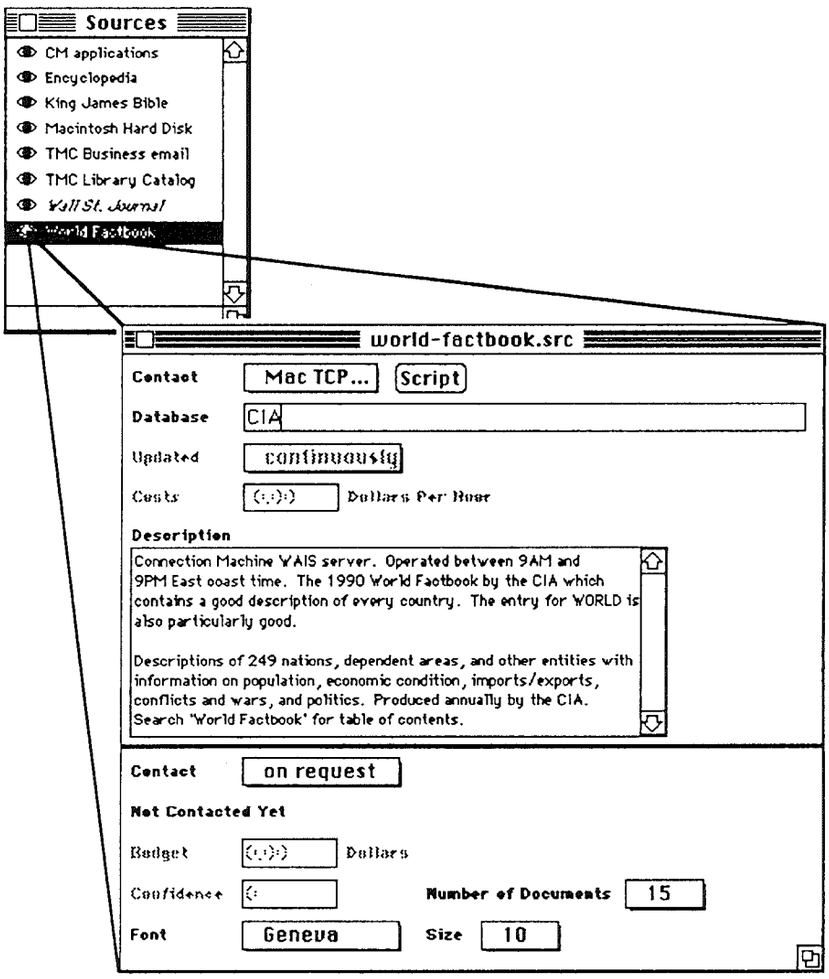

FIG. 5. Double clicking on a source icon opens a source window.

The WAIS interface for the X-Windows environment was developed for the Internet experiment to provide an X-Windows–based interface for a growing community. It was built to look as much like the Macintosh WAIS interface (WAIStation) as possible, given the limitations of the freely distributed X-Windows software. Since the metaphors in XWAIS are nearly the same as those for WAIStation, a user of one system can easily more to the other, without having to learn much additional information. In fact, the underlying data structures are identical to those in WAIStation, so questions can be copied from a Macintosh to a UNIX machine running XWAIS, and used without modification.

XWAIS supports interactive WAIS access, including question entering, source selection, addition of relevant documents, and pieces of documents. Unlike WAIStation, XWAIS retrieves an entire document when requested, instead of just the parts being viewed. We decided this was acceptable, since the underlying networks for X will most likely be fast.

Since XWAIS runs under X-Windows, and was built for the UNIX operating system, it can take advantage of the tools available for these systems to display a wide range of document formats. A simple filter interface is provided in the application (as an X resource) to allow a user to select the tool required for a given type of document; for example, if the document is a Postscript file, xps can be used to view it. This is a feature not available in any of the other user interfaces described here.

In order to distribute this software without restriction, XWAIS uses the freely distributed Athena Widget set including in the X11R4 release from MIT. Although these widgets don't appear as attractive as some others that are available, they can be used to build a useful interface. Some aspects of this interface are restricted by the nature of the widgets available. XWAIS was built using the Xt X Toolkit Intrinsics, and allows a large amount of customization of the appearance of the display using X resources. The application relies heavily on the Xt resource mechanism, and will not run unless these resources are in place. The

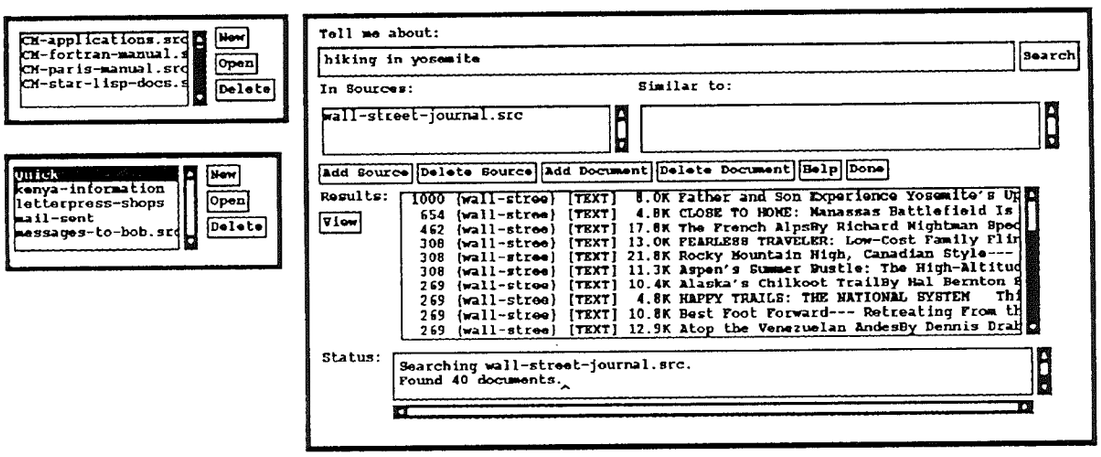

FIG. 6. The XWAIS interface, including the Questions and Sources windows, and an open question.

"object-oriented" feel of these widgets made building the interface rather easy, once the widget with the closest desired functionality was found. Finding the correct widget was the hardest part. Most of the actual behavior of the interface is controlled by "call-backs"—the methods that widgets inherit.

The XWAIS application is actually two separate applications: XWAIS, a simple shell for selecting sources and questions; and xwaisq, the application that actually performs WAIS transactions. The C code in xwaisq is also used in waisq, the shell-support program for GNU Emacs WAIS. This allows users to use simple UNIX facilities to submit questions created by xwaisq using waisq (e.g., a crontab entry to periodically query a server).

The implementation for XWAIS was done in C (6k lines) using the X11R4 release of X-Windows from MIT, the Xt X Toolkit Intrinsics, and the Athena Widget Set, included in the X-Windows release.

XWAIS is a text-based user interface built in a graphical window environment. Some additional graphical metaphors would be desirable, but the limited widget sets precluded that. It would take a considerably larger amount of work to add much graphics to this application. Perhaps some other X toolkit would provide simpler methods for doing this.

GNU Emacs WAIS Interface: GWAIS

GWAIS at a Glance

Target Machine	Terminals on UNIX machines
Effort	2 person-months
Number of Users	500
Status	finished, freely distributed
Language	Gnu-Lisp, and C

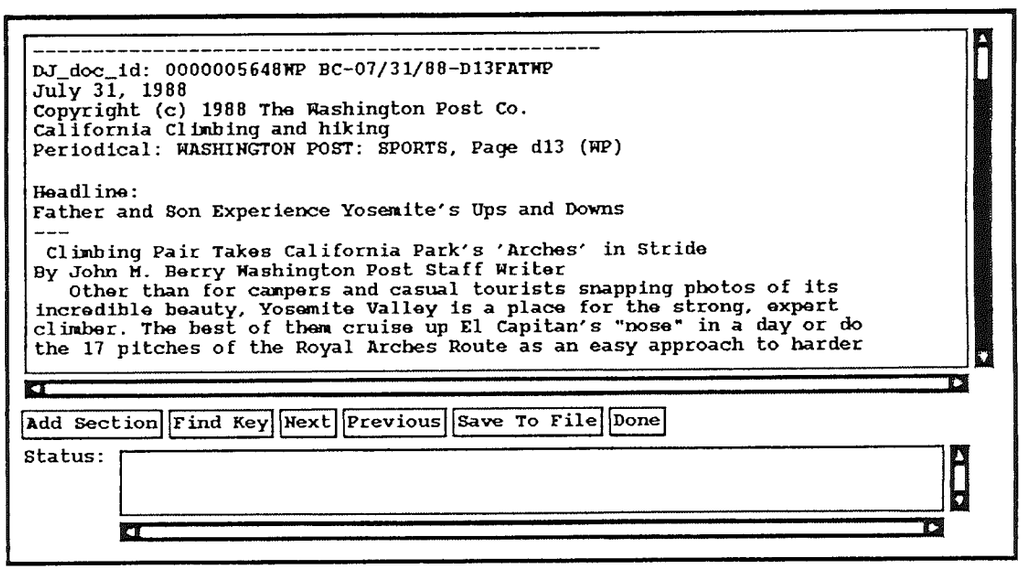

FIG. 7. A document displayed in the XWAIS interface.

Communications	TCP/IP
Designer	Jonathan Goldman
Organization	Thinking Machines
Availability	Available anonymous FTP from /public/wais/wais*.tar.Z@think.com
Design Goals	Copy WAIStation so that we can leverage one design, Use precedent from other GNU-Emacs applications: RMAIL, dired
Used	Used in the Internet experiment with heavy use by some GNU-Emacs users.
Problems	Dealing with the directory of servers. Using passive alerting.

The WAIS interface on GNU-Emacs/UNIX (GNU) was developed specifically for the Internet experiment for a technically strong user population. The reasons it was developed were: the large number of Emacs users, the extensibility, the ubiquitous nature of character display terminals, and the component nature of Emacs, which meant WAIS could be integrated into e-mail, b-boards, and programming tools.

The design of the interface was a cross between WAIStation and other Emacs interfaces. The direct manipulation of WAIStation was replaced by command keys, as is common in Emacs applications. The choice of command keys were modeled on the dired and RMAIL Emacs applications.

GWAIS allow users to access the interactive features of WAIS: question entering, relevance feedback, displaying document, and source selection. An extra feature, not found in the other interfaces, is an interface to an indexer for creating sources, but it appears that this feature is not heavily used. Furthermore, it allows questions to be saved, but it depends on the user to automate the update of questions and sources using cron or other UNIX tools. Graphic documents can be displayed on X-Windows terminals if the user has set up the environment variables.

The implementation of GWAIS was in Emacs Lisp (2K lines) and in C code (3K lines). About half of the time of a typical search and retrieval is spent in reading the data into Lisp.

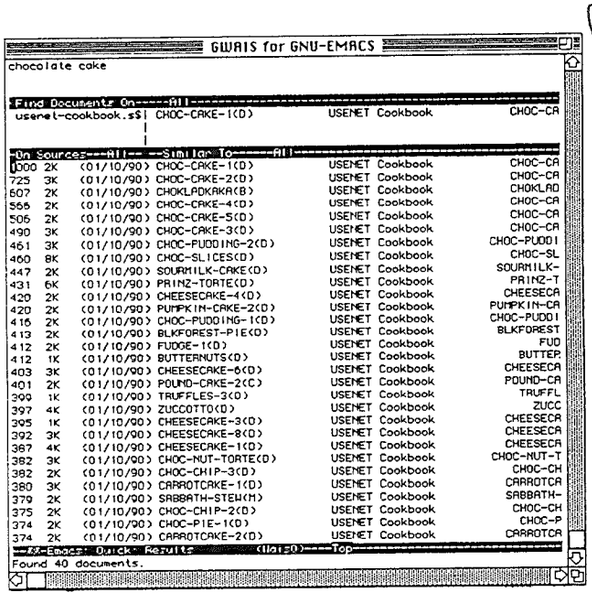

FIG. 8. The GWAIS interface, displaying the results of a relevance feedback search.

Screen-Based (Terminal) WAIS Interface: SWAIS

SWAIS at a Glance	
Target Machine	Terminals connected to UNIX systems
Effort	1 person-month
Number of Users	900
Status	Beta
Language	C
Communications	TCP/IP
Designer	John Curran
Organization	NSF Network Service Center
Availability	To be included in WAIS release, anonymous FTP from /public/wais/wais*.tar.Z@think.com
Design Goals	Highly portable, provide straightforward user interface, utilize existing application key mappings (rn, vi, Emacs), support multiple servers per query, allow for personal "source" directory and a common source directory, allow for useful source discovery via searches, provide simple active tool with little state (no question storage, relevance feedback, or passive notification).
Used	Internet users via Telnet: K-12 students, educators, user services staff, librarians, and (occasionally) network staff.
Problems	Dealing with the directory of severs. Lack of information in many server-returned records. Providing simple and uniform nomenclature. Planning for large numbers of sources.

To open WAIS to a wider community of users, an interface was developed to run on dumb terminals or over Telnet sessions. It is called "SWAIS" for Screen WAIS since it uses a character display terminal screen for the interface. The user communities that this interface can serve are dial-in users, Telnet users, and low-end terminal users.

The design of the interface involved three screens: a single screen listing all known servers that the user could pick from, a list of search result documents headlines, and a document display screen. Listing all servers and allowing users to select which servers to use encourages users to ask questions of multiple servers. Unlike the other interfaces, the sources list shows what site runs it and how much it costs (if anything). The resulting document screen includes headlines and number of lines, but its innovation is to show the source.

It does not handle relevance feedback or downloading new sources from the directory of servers. Another drawback is using it with large numbers of sources, since moving around the list requires scrolling. On the other hand, this server has proven to be very popular on the Internet.

Ed: I have requested paragraphs in single space from author (col 1)

Au: Should these be LISP? (not changed as marked in my original request)

Au: Is this word (dired) correct? Not answered in my original query. (col 1)

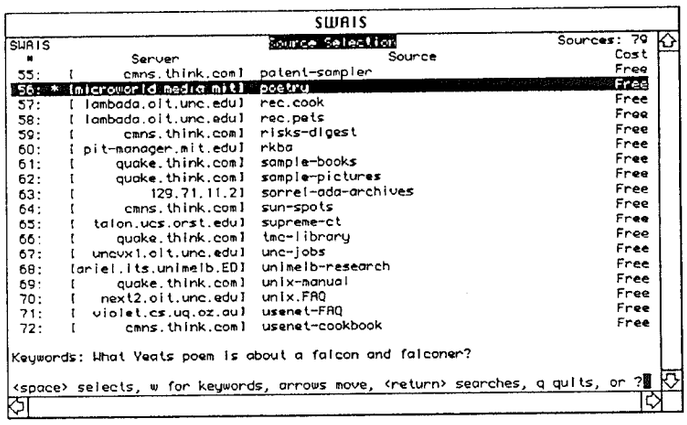

FIG. 9. The SWAIS query building screen. The poetry source is selected, and search terms are entered. This interface does not currently support relevance feedback.

Au: Clear? (col 1) YES

Ed: *Note* I've requested single space c in paragraphs from author. (col 1&2)

Because of its ease of use, all a user has to do is use Telnet to a specific machine to use it.

The Rosebud Interface: Reporters and Newspapers on the Macintosh

Rosebud at a Glance

Target Machine	Macintosh II, color screen
Number of Users	25
Status	Finished; internal use
Language	Smalltalk, MPW-C
Communications	TCP/IP using IPC package
Designers	Charlie Bedard, David Casseres, Steve Cisler, Tom Erickson, Ruth Ritter, Eric Roth, Gitta Salomon, Kevin Tiene, Janet Vratny-Watts.
Organization	Apple Computer
Availability	Only internally to Apple ATG
Design Goals	Serve as research platform for interface and architectural explorations. Allow ordinary users to create personalized information flows; support passive alerting, scanning, and capture of information.
Used	Used in various internal tests; not available for the Internet experiment.
Problems	No good interface mechanisms for providing users with convenient access to large numbers of servers.

Rosebud is a project within Apple Computer's Advanced Technology Group. Its principle objective is to serve as a platform for investigations into what is needed to make remote information accessible and useful to ordinary Macintosh users. The investigations have two foci: human interface components and techniques, and system architecture issues. In this article, we focus exclusively on the human interface aspects of Rosebud.

The Rosebud Server System is similar to the WAIS system in that it uses the Z39.50 protocol to access multiple, remote databases; it differs from them in that it contains extra underpinnings for making information access an internal part of the Macintosh environment. Specifically, the Rosebud Server System allows users to create autonomous, ongoing "agent" processes which access, update,

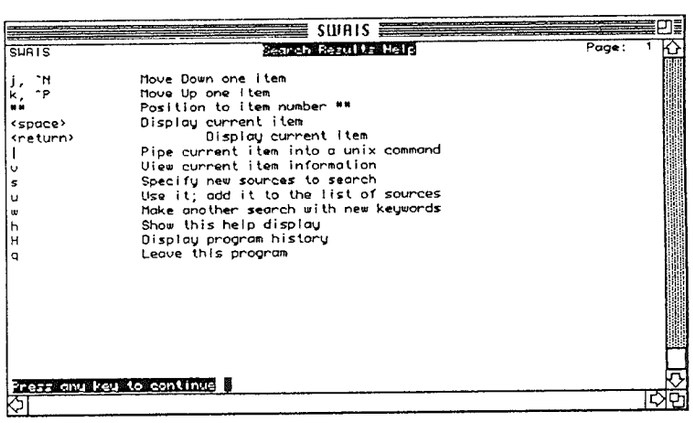

FIG. 10. The SWAIS help screen.

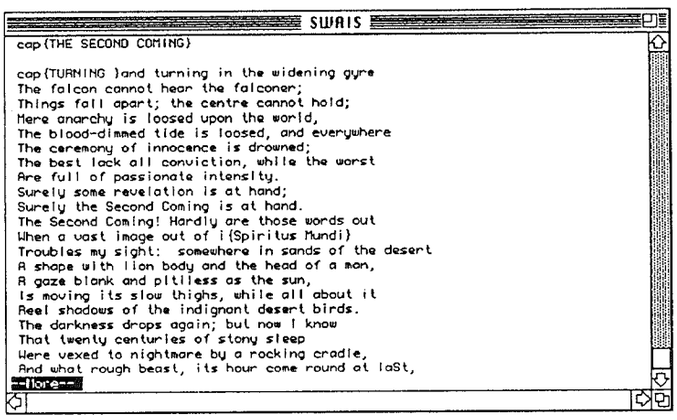

FIG. 11. A document displayed in SWAIS.

and present information from local and remote sources. The Rosebud system does not currently provide access to the Internet WAIS servers (for reasons of network security, rather than basic incompatibilities), and it is not publicly available.

Design Rationale

The design of the Rosebud interface began with a study of the practices and problems of ordinary information users. The principle focus was on information users at KPMG Peat Marwick in San Jose, the original client site for WAIS; in addition, several groups of users of online information services within Apple were also studied (Erickson & Salomon, 1991). Interviews with accountants at Peat Marwick enabled the designers to put together a schematic of how information (mostly paper-based information) flowed through their offices (Fig. 12).

Several features of this schematic informed the design of Rosebud. First, information typically came to the accountants via newspapers, magazines, and memos; instances where the accountants went out of their way to search for information were less frequent. Second, the accountants never talked about "reading" information; they always spoke of scanning, or skimming it—they did not have time to read it. This suggested that a good interface should provide a way for the users to scan retrieved information quickly. Third, accountants remarked that they discarded most information, including information that might be useful. Potentially useful information was discarded for two reasons: the accountants did not have the physical space to store everything, and they knew from experience that if they tried to save too much, they would not be able to find anything later when they actually needed it. This suggested that giving users access to remote information was just half the problem; users also needed tools for archiving, organizing, and reretrieving information. Finally, when users did come across information that seemed worth saving, they typically would cut it out (the accountants used, almost exclusively, paper-based information), and then they would annotate it by circling, underlining, or jotting a few notes in the margin. Annotation turned out to be an important concept: not only did it help the user who annotated when the information was reretrieved later, but it also helped others scan the information more quickly when copies were passed on to them.

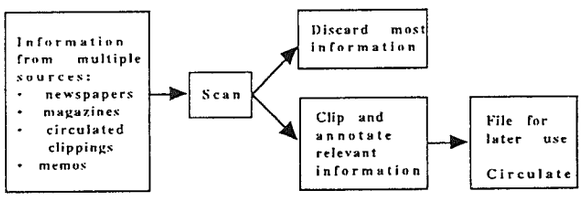

FIG. 12. Information flow through accountants' offices.

The consequence of these observations was a design for a system which allowed users to define topics of interest which would be retrieved automatically, and would then permit them to scan those items and save them into an environment where they could be annotated, organized, and reretrieved.

Human Interface

The Rosebud interface design has three components: reporters, newspapers, and notebooks. Reporters are for retrieving information. Users give reporters assignments which specify what to look for, and where to look. This is shown in Figure 13: users enter words describing the information in which they are interested, check off the information sources they wish the reporter to search, and, if they so choose, automate the reporter so that it searches the databases on a daily or weekly basis. Upon pressing the "Search" button in the assignment window, a reporter is created, performs the search, and returns with a list of results (Fig. 14).

The reporter window (Fig. 14) provides users with a variety of ways to look over their results and refine their queries. The results are shown in the "Best Guesses"

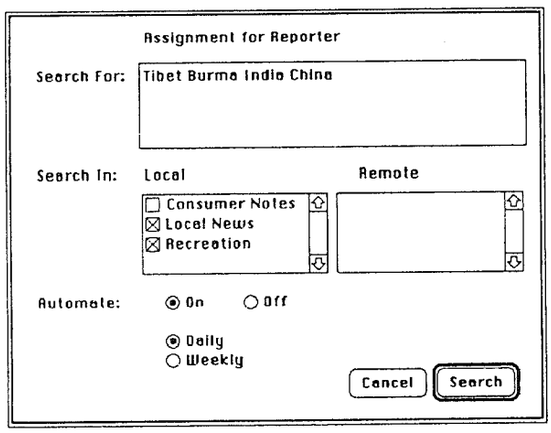

FIG. 13. Creating a reporter—the assignment window.

pane. (The name "Best Guesses" was chosen to provide some indication that inaccuracy could be expected; our observations of users had shown that they were often mystified by some of the items that showed up as the results of searches.) The asterisks to the left of items indicate their relative relevance, and the pop-up menu above the pane allows users to order the list by date or relevance. Simply selecting an item shows a preview of it—a short excerpt with search terms highlighted in color and boldface (Fig. 15). Previews are useful because users can get a look at a small part of the item without incurring the overhead of downloading the whole article over the network. Users also have the options of saving articles to their disks or opening them for viewing. Finally, having looked over their results, users can refine their search in the bottom pane of the window.

The above sequence occurs whenever a user creates a new reporter. However, since users are likely to use many

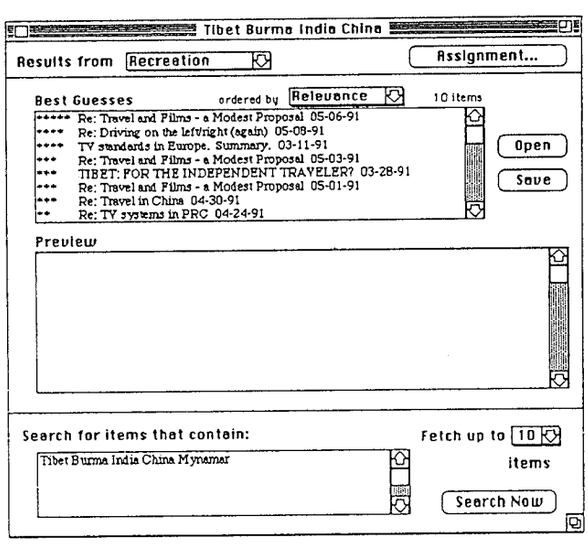

FIG. 14. The reporter window contains the results of the search and provides means for previewing, opening, and saving results.

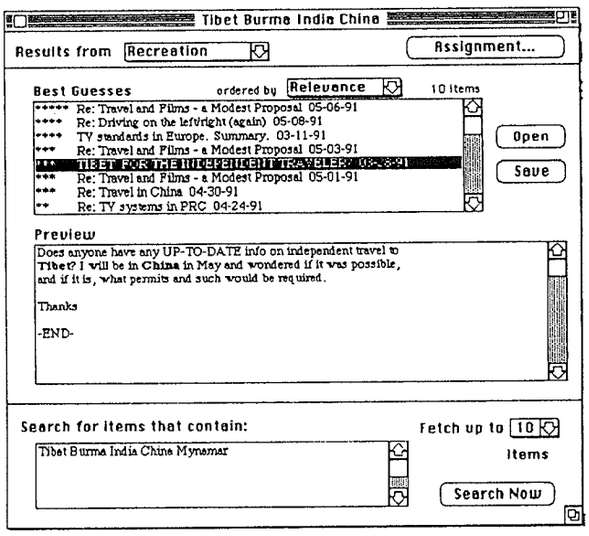

FIG. 15. The reporter window makes it easy to scan through hits. Clicking on a retrieved item generates a preview which shows an excerpt to the hit with the search terms (Tibet and China) highlighted in boldface. The user can refine the query in the lower pane of the window.

reporters, and because the initial user studies indicated that ways of skimming through incoming information were important to the accountants, the newspaper was provided to support rapid scanning of new information. The model of a newspaper is quite simple (Fig. 16): on the left is an index column which contains the names of all reporters, and to the right are two columns of news. Each reporter "owns" one news column and publishes the title, date, and an excerpt of each item in its column. The columns scroll independently, using "minimalist" scroll bars to prevent the multiple scroll bars from visually overloading the screen. If an excerpt seems interesting, double clicking on it opens the full article in a window, from which it can be viewed, printed, or saved. Thus, rather than having to open up a dozen reporters every morning to see what's new, the user can go to one place, the newspaper.

The newspaper can also serve as a control center for the Rosebud interface. The user can open a reporter by clicking on its name or icon at the top of its new column. Consequently, if a reporter's column has strayed from the desired topic, the user can quickly get to the reporter and revise its assignment. The index also lists inactive reporters (those either not automated, or that have not found anything new since the last newspaper), so they also can be opened, automated, or otherwise adjusted.

A third component of the Rosebud interface—the notebook—was designed but not implemented. Notebooks are environments within which users may save, annotate, and organize retrieved information. Notebooks were designed in response to the observations of Peat Marwick accountants, which indicated the need for an environment which supported the way accountants worked—in particular, notebooks were intended to support annotation and refinding

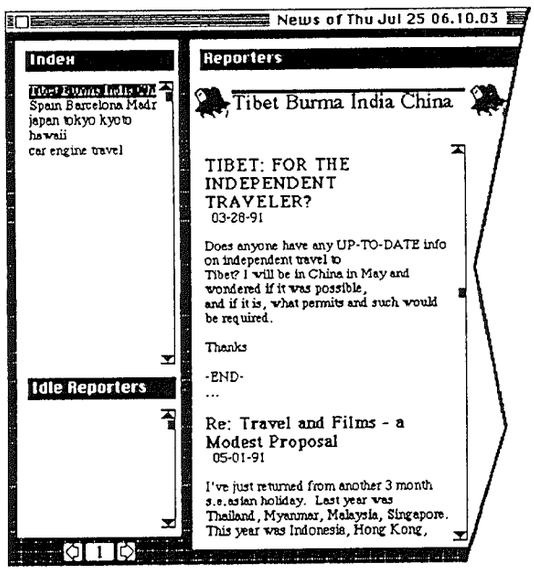

FIG. 16. The newspaper allows users to quickly scan through new items retrieved by the reporters which are working automatically.

retrieved information at a later date. A particularly nice feature of the notebook design was its use of annotations as landmarks for refinding information. The notebook design, and its rationale, is described in Erickson and Salomon (1991).

Implementation

The Rosebud system consists of six parts: (1) a human interface application written in SmallTalk/V (to facilitate the rapid changes in the interface necessary to effectively conduct interface design research); (2) a search manager package which implements the autonomous agent functionality and formulates Z39.50 queries for (3) remote Z39.50 servers implemented in MPW C that automatically index items placed in their input folders by (4) HyperCard stacks that download new items from a Net News server; (5) a file manager component (MPW C) that does all of the file I/O and compaction for reporters and newspapers; and (6) directory servers which allow the various components to find one another. All of these components are written as separate applications and communicate with one another using a prototype IPC that runs over TCP/IP. The file manager and search manager applications run in the background under MultiFinder, enabling Rosebud to access information and construct newspapers while the human interface application is not running. Like the other WAIS interfaces, Rosebud uses the WAIS protocol package. The human interface was designed for Macintosh II class machines, with 13-inch color screens.

Observations and Testing Results

The Rosebud human interface was subjected to informal testing on 14 users. Users were told only that Rosebud was an application for finding information, and then given a particular topic to find information on. They were given no help or documentation. Note that, although informal, this type of testing is very stringent in that users approach the application knowing almost nothing about what it is, or why they would actually use it. Data collection consisted simply of recording their questions, observations, and problems as they went along, administering a posttest questionnaire, and then asking them a few, open-ended questions. Here are a few of the more general observations.

- Over 80% of those who tried the Rosebud interface responded very positively to it, and said that they would use something with its capacities as part of their daily work routine. Two thirds of users indicated that they would usually use newspapers to browse through information (instead of reporters).
- At the end of the test, over two thirds of the users said they liked the metaphors of reporters and newspapers; however, almost all users had some difficulty in getting started. The typical problem was that users did not associate reporters with a way of retrieving information. When asked to find information, users first looked for an item called "search"; when they did not find this, they usually turned to the newspaper, which is, in fact, where they look for information on a daily basis. It is possible that this problem can be remedied by minor interface changes (e.g., putting a "New Reporter" item in a search menu); alternatively, it may be that the metaphor is inappropriate.
- A number of users were led astray because they had conceptual models of information retrieval based on their familiarity with query languages and structured databases. Such users tended to be wary of entering search terms because they were not sure of the appropriate syntax, and did not understand what "relevance" meant. Those who did know the meaning of relevance wanted to know how the information server calculated it.
- Users liked previews—especially the feature of highlighting keywords in boldface. They wanted to see boldface keywords in the newspaper and article windows. Users also wanted the ability to select text in the newspaper and article windows and change the style or font themselves, so that they could annotate significant items. This parallels practices observed in our initial observations of accountants, where we found that annotation plays several important roles.
- A variety of low-level interface problems, due to terminology or graphic design were discovered. Some examples: users did not usually recognize the asterisks in the "Best Guesses" window as indicators of relevance; users did not think that "idle reporters" was a good name, and said that it was very important to distinguish between reporters that had found nothing, and those that were not looking.

The testing described above focused on how usable Rosebud was with users' first exposure. In the next phase of testing, a small set of users will be observed over the course of a month, in which they have the option of using Rosebud from their desktop machines to access meaningful data. This phase of testing will allow a more realistic assessment

of Rosebud, in that it will last long enough to permit users to build up their own set of reporters, and to access newspapers that contain information of personal import.

Conclusion

This article has described five interfaces developed to provide access to distributed systems of information servers. The interfaces presented here were developed with different constraints in mind, so it is not useful to compare them directly; instead, they may serve as examples of differing responses to issues such as screen size, workstation power and intelligence, communication speeds, and user needs and practices.

The interfaces designed so far have addressed some of the critical issues for end-users to accomplish interactive searches in a wide-area network. These include ways of finding which information servers contain relevant information; supporting searching by ordinary users; and supporting browsing of, and passive alerting about, newly retrieved information. The alerting aspects of the interfaces have not been tested much in this environment due to the lack of appropriate data sources for this type of searching. It is probably fair to say that any of the design solutions described here can be improved upon by further work.

The WAIS Internet experiment has revealed a number of issues requiring further work. In the Internet environment, we have observed (in the logs of user queries) that users have a difficult time finding out what is in a database, thus demonstrating there is a lack of browsing or scanning facilities in the interfaces, protocol, and servers, as well as a general shortage of descriptive information about databases.

Finally, a variety of other issues were raised during the studies of the Peat Marwick accountants who have received little or no work. Document layout is one such problem. Accountants mentioned that sometimes they want to retrieve documents not because of the information they contain, but to look at their layouts (accountants will often examine successful proposals to a client when preparing a new proposal). More generally, users regard pictures, diagrams, tables, and charts as essential components of a document's content. Unfortunately, support for different document formats, and for the retrieval and display of nontextual information within them, is very limited on most existing clients.

Another issue is called the boilerplate problem. Accounting documents often contain a large amount of boilerplate—standard text which varies little from document to document. What tools are needed to allow users effectively to retrieve, order, and browse a large set of documents which are 95% similar? Note that boilerplate is characteristic of a wide variety of business proposals and legal documents, not just accounting documents. In fact, the analog to boilerplate occurs in scientific documents in which standard terms and descriptions are used to describe procedures and methods used in an investigation.

A number of other issues remain to be addressed. Users are very interested in being able to see what queries other users are conducting, and what information servers and articles are most popular. A frequent suggestion is to allow users to rate the "goodness" of articles they retrieve. However, in a commercial setting, information about the kind of questions being posed by a particular company or person can be revealing and valuable. Clearly, the utility that such information could provide must be balanced by concerns about confidentiality and privacy, and mechanisms for user control of descriptive information are essential. Other issues include how to control the pricing, copyright, and distribution issues which accompany "for-pay" information.

In summary, there is an immense amount of work to be done. A central part of this work involves further research and development of interfaces. We have made the WAIS system publicly available in the hope that designers will find that it—with its common protocol and defined infrastructure—can serve as a platform from which to pursue these and other research issues.

Acknowledgments

At Thinking Machines: Ottavia Bassetti, Franklin Davis, Patrick Bray, Danny Hillis, Rob Jones, Barbara Lincoln, Gordon Linoff, Chris Madsen, Gary Rancourt, Sandy Raymond, Tracy Shen, Craig Stanfill, Steve Schwartz, Robert Thau, Ephraim Vishniac, David Waltz, and Uri Wilensky.

At the National Science Foundation: Suzi, Alex, et al.

Appendix

The success of a distributed system of information servers depends on a critical mass of users and information services. To encourage development and use, Thinking Machines is making the source code for a WAIS protocol implementation freely available. While this software is available at no cost, it comes with no support. We hope that it will facilitate others in developing servers and clients.

For more information, please contact:

Barbara Lincoln Brooks (barbara@wais.com)
WAIS Inc.
1040 Noel Drive
Menlo Park, CA 94025

References

Davis, F., Kahle, B., Morris, H., Salem, J., Shen, T., Wang, R., Sui, J., & Grinbaum, M. (1990). *WAIS Interface Protocol functional specification,* Thinking Machines Corporation, April. Available via anonymous ftp: /pub/wais/doc/wais-concepts.txt@quake.think.com or wais server wais-docs.src

Dongarra, J., & Grosse, E. (1987). Distribution of mathematical software via electronic mail. *CACM, 30,* pp. 403–407.

Egan, D., Remde, J., Gomez, L., Landauer, T., Eberhardt, J., & Lochbaum, C. (1989, January). Formative Design-evaluation of SuperBook. *ACM Transactions on Office Information Systems.*

Erickson, T., & Salomon, G. (1991). Designing a desktop information system: Observations and issues. In *Proceedings of the ACM Human Computer Interaction Conference* (pp. 49–54). New York: ACM.

Au: Please supply page nos. & vol. (Egan et al.) Au: City okay? YES (Erickson & Salomon)

January 1989, Volume 7, No. 1, pp. 30-57.

Ginther-Webster, K. (1990). Project mercury. *AI Magazine,* pp. 25–26.

For information on the Free Software Foundation and the GNU project, see: /pub/gnu/*@prep.ai.mit.edu. For Emacs see: /pub/gnu/emacs-*.tar.Z@prep.ai.mit.edu. Free Software Foundation, 675 Massachusetts Avenue, Cambridge, MA 02139.

Kahle, B., & Medlar, A. (1991). An information system for corporate users: Wide area information servers. ~~Apple~~ *Online Magazine,* September, pp 56-60.

Kahle, B., Goldman, J., Morris, H., & Sten, T. (1991). Electronic publishing experiment on the internet—wide area information servers.

Malone, T., Grant, K., & Turback, F. (1986). The information lens: An intelligent system for information sharing in organizations; In human factors in computing systems. In *CHI 86 Conference Proceedings* (pp. 1–8) Boston, MA, New York: ACM.

MOSIS - for more information on the mosis fabrication facility, write to mosis@mosis.edu.

For more information on the Mosis fabrication facility write to mosis@mosis.edu

NeXT Computer Inc., 900 Chesapeake, Redwood City, CA 94063 (415) 366-0900.

Z39.50-1988: Information retrieval service definition and protocol specification for library applications, National Information Standards Organization (Z39), P.O. Box 1056, Bethesda, MD 20817 (301) 975-2814. Available from Document Center, Belmont, CA (415) 591-7600.

ON Technology Inc., 155 Second Street, Cambridge, MA 02141 (617) 876-0900.

Verity Inc., 1550 Plymouth Street, Mountain View, CA 94043 (415) 960-7600.

Au: Which month is it (April or September)? Please given page nos. (Kahle & Medlar)

Au: Please give page nos., journal or book? (Kahle, et al.)

This was not published

MOSIS should be an entry in references, with NEXT being next item.